

*Wikimedia Foundation v. NSA*  
No. 15-cv-0062-TSE (D. Md.)

# Plaintiff's Exhibit 5

**IN THE UNITED STATES DISTRICT COURT  
FOR THE DISTRICT OF MARYLAND**

WIKIMEDIA FOUNDATION,

*Plaintiff,*

vs.

NATIONAL SECURITY AGENCY /  
CENTRAL SECURITY SERVICE, *et al.*,

*Defendants.*

No. 1:15-cv-00662-TSE

**DECLARATION OF TILMAN BAYER**

I, Tilman Bayer, declare:

1. I am a resident of San Francisco, California, over the age of eighteen. I have personal knowledge of the facts stated in this declaration and if called to testify I could and would testify competently thereto. I am providing this declaration in my capacity as an employee of the Wikimedia Foundation, Inc. (“Wikimedia”).

2. I am a Senior Analyst in Wikimedia’s Product Analytics team, and have been a full-time employee of the organization since 2012. My responsibilities include the reporting of pageview statistics and other key web traffic metrics to Wikimedia’s executives and board. I hold degrees in mathematics from the University of Cambridge (Certificate of Advanced Study in Mathematics) and the University of Bonn (diploma, equivalent to a Master’s degree in the US).

3. As explained in further detail in the Declaration of Michelle Paulson (Exhibit 3), Wikimedia is a 501(c)(3) nonprofit organization dedicated to encouraging the growth, development, and distribution of multilingual educational content, and to providing the full content

of these “wiki”-based projects to the public free of charge.<sup>1</sup> Wikimedia operates twelve free-knowledge “Projects” on the Internet, including Wikipedia, a free-access, free content encyclopedia that is the Internet’s largest and most popular reference work, and one of the top ten most-visited websites in the world. At the time of the filing of the Amended Complaint in 2015, Wikipedia received more than 400 million visitors each month. *See* Exhibit 44 (WIKI0008271). As of February 2018, the site has grown to contain more than 47 million articles in over 288 languages, and in 2017 it received visits from more than 1 billion unique devices each month.

4. Wikimedia performs a variety of activities that support the Wikimedia movement. Most relevant to my testimony here, Wikimedia provides technical infrastructure for the Projects, and operates over 1000 servers located in the United States, in the Netherlands, and since March 2018, in Singapore. In response to Defendants’ written discovery requests in this litigation, I was tasked with quantifying the volume of Internet communications that are transmitted by Wikimedia’s servers. My analysis quantified communications by overarching category, Internet protocol, and country of origin. The results of these analyses were produced during discovery. *See* Exhibits 14, 42, and 43. I also quantified and verified other statistics relevant to the volume of Wikimedia’s communications, a subset of which are produced herein. *See infra* ¶ 29.

## **I. TECHNICAL BACKGROUND**

### **A. HTTP/S Requests and Responses**

5. A brief description of how computers communicate over the Internet is helpful to understanding the nature of Wikimedia’s communications with its users. More fulsome explanations of these technical processes are provided in the expert Declaration of Scott Bradner (Exhibit 1).

---

<sup>1</sup> A “wiki” is an application that allows collaborative modification, extension, or deletion of its content or structure.

6. The hypertext transfer protocol, or HTTP, is the foundation for data communication over the world wide web and facilitates the exchange of information between computers. HTTP is a “request/response” “Application Layer” protocol that governs how communications occur between “clients” and “servers.” A “client” is often a software application, such as a web browser, that initiates a request to connect with a “server,” which may be a computer hosting a website. To access a particular webpage, a client must send at least one HTTP request to the relevant server. The number of requests required for an Internet user to access a particular webpage depends on the number of graphics, videos, and other specialized components featured on the page. A server responds to a request by providing the requested “resource.” Uniform Resource Locators, or URLs, both identify a resource and describe its location or address. When an Internet user enters a URL address into their web browser using the “http” or “https” web address format, or when a user clicks on a “link” to a webpage, they are actually instructing their web browser (the client) about which resource to request and where to find it.

**B. HTTPS and Wikimedia Projects**

7. Hyper Text Transfer Protocol Secure (HTTPS) is the secure version of HTTP. Web browsers use the URL convention of replacing “http:” with “https:” to indicate that the browser is to communicate with the web server using the “Transport Layer Security” (TLS) protocol instead of the unencrypted protocol used by HTTP. HTTPS operates between the application and transport layers, and encrypts the HTTP application layer communication.

8. Over the past several years, Wikimedia has transitioned to the use of HTTPS by default. *See* Paulson Decl. ¶¶ 49-50. Currently, HTTPS is enabled on all Wikimedia Project websites. If a user requests a Wikimedia site using “http:” in their web browser, they receive a “redirect” response from Wikimedia’s servers, which includes the URL information from the user’s request.

**C. Other Internet Communications Protocols**

9. The Transmission Control Protocol (TCP) is a transport layer protocol used to provide a reliable data stream between network nodes. It is used by most major Internet applications including email and the world wide web.

10. The User Datagram Protocol (UDP) is a transport layer protocol that provides a way to send packets from one network node to another network node. Many applications use UDP for transport including some voice and video streaming applications.

11. The Internet Control Message Protocol (ICMP) is a supporting communications protocol used by network devices, including routers, to send operational information, including error messages.

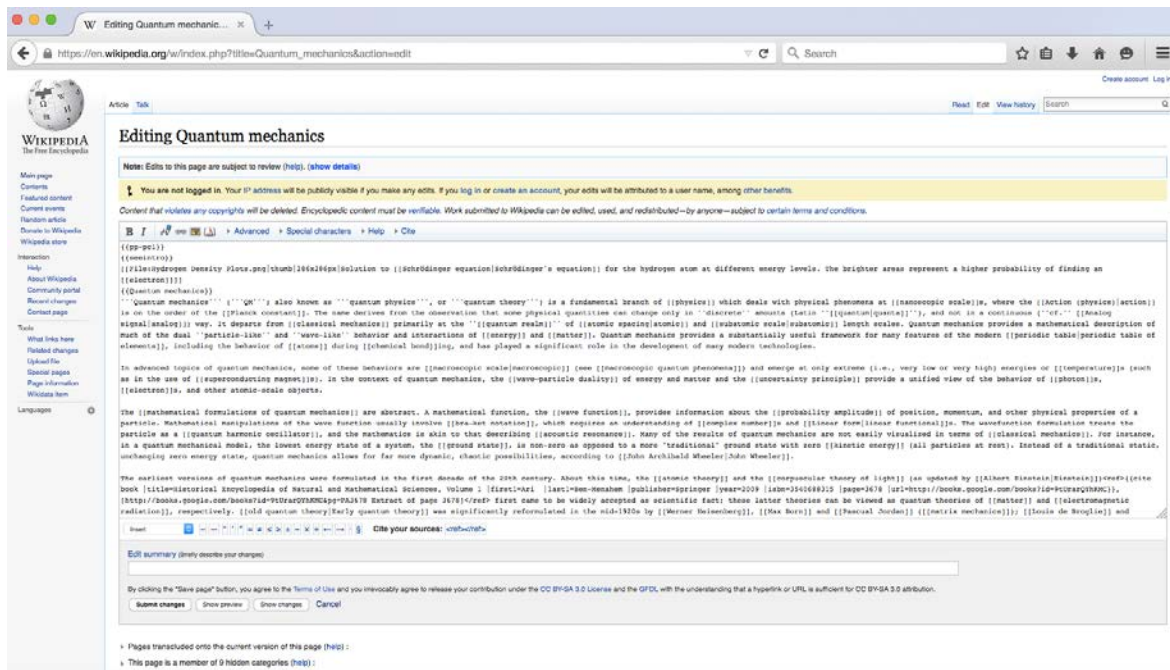
**II. WIKIMEDIA PROJECT PAGES AND FEATURES**

**A. Editing Project Pages**

12. All twelve of the Wikimedia Projects are written, edited, and curated collaboratively online by volunteers. Wikimedia users can contribute to Project websites anonymously, under a pseudonym, or under a username that reflects their real identity. As of March 2018, there have been approximately 3.4 billion edits over the lifespan of the Wikimedia Projects.

13. Users can edit a specific Wikimedia Project page via an “Edit” page, which contains a text box with an editable version of the article. An example of an Edit page is depicted in Figure 1 below.

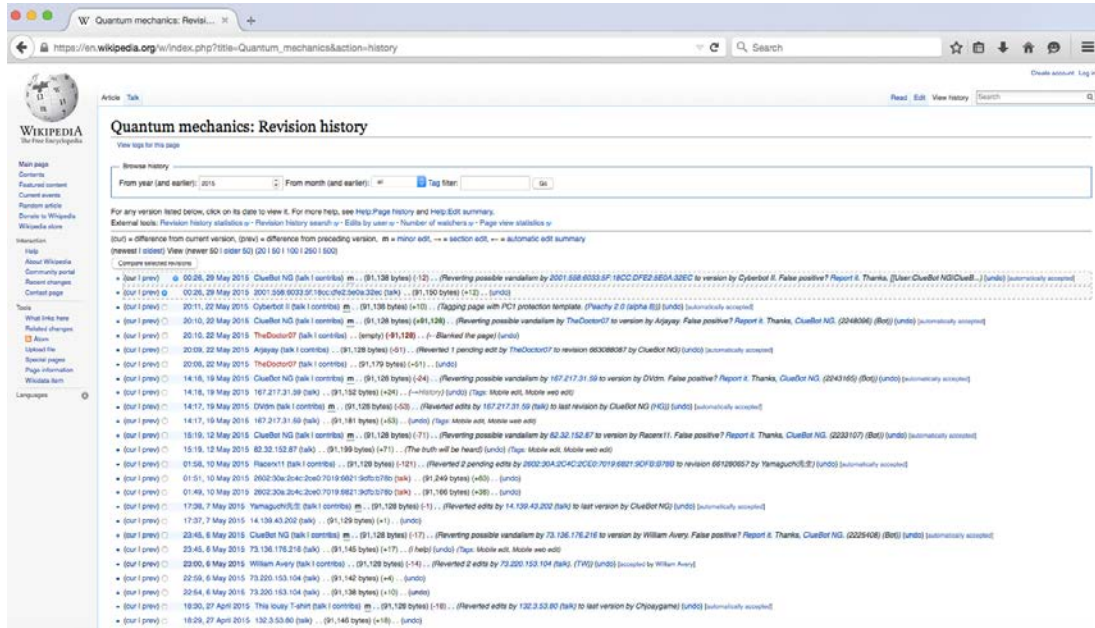
Figure 1



14. Any revisions on this edit page will generally remain private until the user selects the “Publish changes” option (earlier named “Save page” or “submit changes”) at the bottom of the screen. If a user selects “show preview” while editing a draft, the changes are communicated between the user’s browser and Wikimedia’s servers. After the user saves her changes, the user’s revisions are publicly visible on the page.

15. The revision history for every page on each of the twelve Projects is publicly visible by clicking on the “View history” or “History” tab on the right side of the screen. This revision history includes unique edits made by individual users. An example of a revision history page is depicted in Figure 2 below.

Figure 2



16. Figure 3 displays a portion of the revision history of the Wikipedia page on “Quantum mechanics.” As illustrated in this Figure, when a user who is not logged in as a registered user edits the article, the user’s IP address information is visible next to the contributions from that IP address. When a user contributes while logged in to his account, the corresponding revision history page does not reveal the user’s IP address. Instead, the revision history page displays the user’s username; a link to the user’s contributions; and a link to his “User page.” That history of the user’s contributions and edits is also available via a “Special: Contributions” link from the User page.


Figure 3

- (cur | prev) ○ 20:08, 22 May 2015 [TheDoctor07](#) (talk | contribs) .. (91,179 bytes) (+51) .. (undo)
- (cur | prev) ○ 14:18, 19 May 2015 [ClueBot NG](#) (talk | contribs) **m** .. (91,128 bytes) (-24) .. (Reverting possible [automatically accepted])
- (cur | prev) ○ 14:18, 19 May 2015 [167.217.31.59](#) (talk) .. (91,152 bytes) (+24) .. (→History) (undo) (Tags: Mobile edit, Mobile web edit)
- (cur | prev) ○ 14:17, 19 May 2015 [DVdm](#) (talk | contribs) **m** .. (91,128 bytes) (-53) .. (Reverted edits by 167.217.31.59 to last revision by ClueBot NG (talk) [undo])
- (cur | prev) ○ 14:17, 19 May 2015 [167.217.31.59](#) (talk) .. (91,181 bytes) (+53) .. (undo) (Tags: Mobile edit, Mobile web edit)
- (cur | prev) ○ 23:00, 6 May 2015 [William Avery](#) (talk | contribs) .. (91,128 bytes) (-14) .. (Reverted 2 edits by 73.200.153.104 (talk) (talk) [undo] [auto] [summary])
- (cur | prev) ○ 22:59, 6 May 2015 [ClueBot NG](#) (talk | contribs) **m** .. (91,142 bytes) (+4) .. (undo)
- (cur | prev) ○ 22:56, 6 May 2015 [73.200.153.104](#) (talk) .. (91,138 bytes) (+10) .. (undo)
- (cur | prev) ○ 18:20, 27 April 2015 [This is my 7-ohh talk | contribs](#) **m** .. (91,128 bytes) (-10) .. (Reverted edits by 132.3.53.80 (talk) to last version by Chrysogames [undo] [auto] [summary])
- (cur | prev) ○ 18:20, 27 April 2015 [132.3.53.80](#) (talk) .. (91,148 bytes) (+10) .. (undo)

**B. Searching and Navigating Wikipedia**

17. As shown in Figure 4, users may search Wikipedia by entering a search term; on the desktop version of the site, shown here, the search box is located in the upper-right corner of the screen.

**Figure 4**



18. Where the search identifies a Wikipedia page with a title that corresponds exactly to the search term, the user is sent directly to that page. The URL of the resulting page incorporates the relevant topic and reveals information about the content of the requested page. For example, when a user employs a search for “substance abuse,” the resulting webpage URL is [https://en.wikipedia.org/wiki/Substance\\_abuse](https://en.wikipedia.org/wiki/Substance_abuse).

19. When a search does not identify a Wikipedia page with a title that corresponds exactly to the search term, the user is sent directly to a “Search results” page. The URL of the search results page incorporates the search term employed by that user. For example, if a user searches for “Kansas City Monsoon”—a fictional event with no corresponding Wikipedia page—he would be directed to a Search results page with the following URL: <https://en.wikipedia.org/w/index.php?search=Kansas+City+Monsoon&title=Special%3ASearch&go=Go>.

20. Additionally, Wikipedia pages include links to other Wikipedia pages on related topics that enable the reader to access those pages directly from the page being viewed. For instance, the Wikipedia page on “Substance abuse” references “anti-social behavior,” and includes



a link to the Wikipedia page on this topic, [http://en.wikipedia.org/wiki/Anti-social\\_behaviour](http://en.wikipedia.org/wiki/Anti-social_behaviour).

21. In 2017, Wikimedia sites received over 237 billion page views, with approximately 74 billion views originating from users in the United States.

### C. User Discussions

22. Wikimedia also engages in communications that permit its users to interact with one another more directly. Wikimedia Project pages feature “Talk” pages (also known as “Discussion” pages) in which users can publicly discuss potential changes to an article or other Wikimedia Project page. When viewing a Wikimedia Project page on the desktop version of the site, the link to the corresponding Talk page appears on the top left of the screen. A depiction of a Talk page tab and an article Talk page is shown in Figures 5 and 6 below.

**Figure 5**

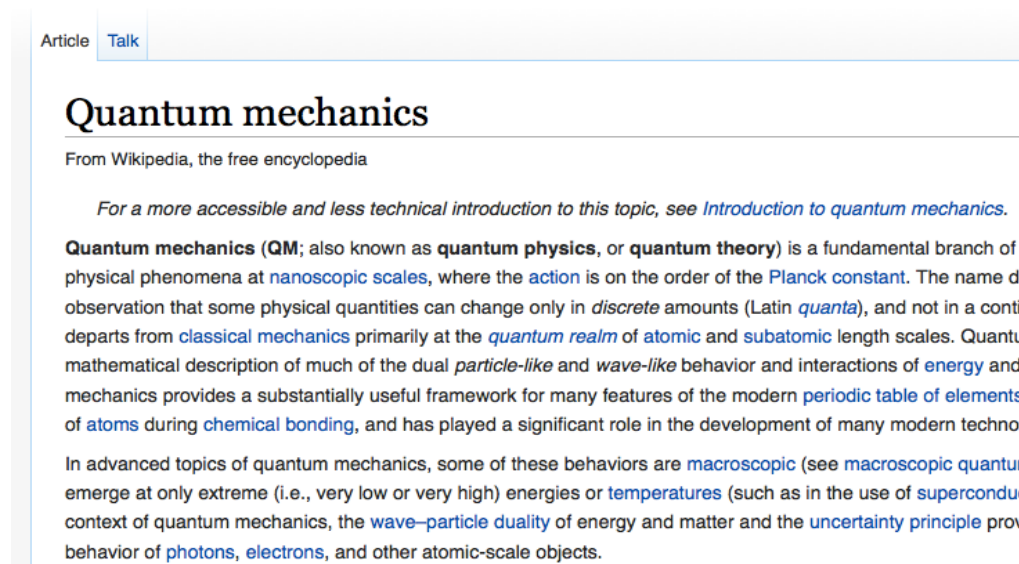
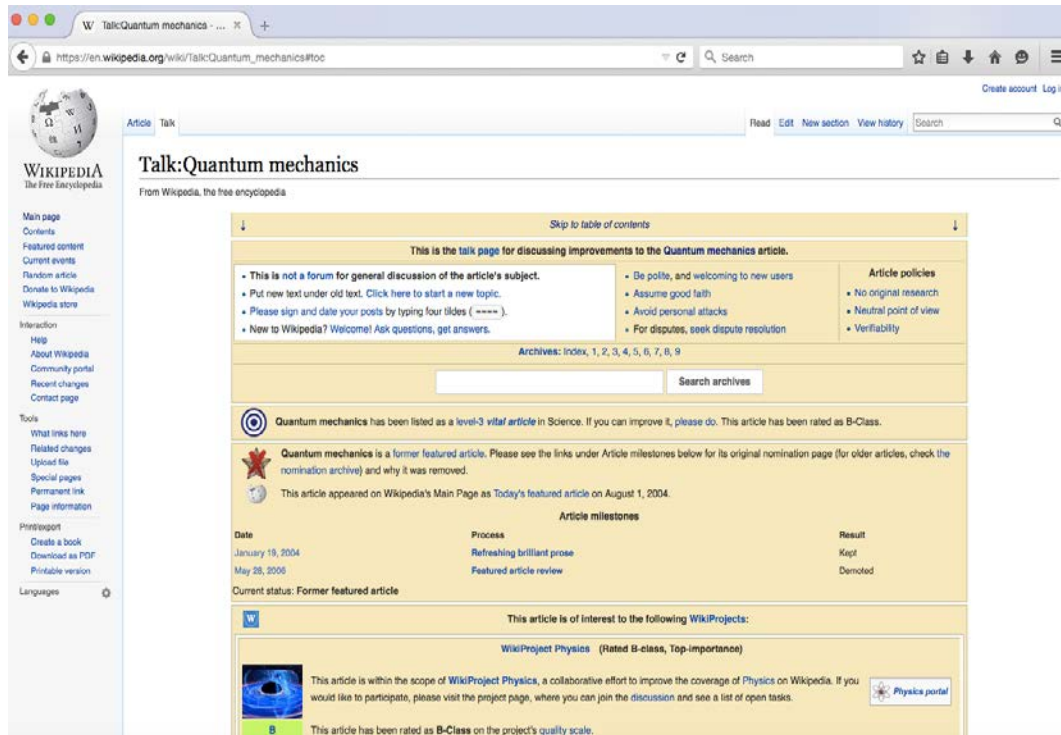


Figure 6



23. User pages also have associated Talk pages. One user may contact another by leaving a message on his Talk page. As with edits, Talk page comments are publicly visible and associated with the commenter's username, or, if the user is commenting without logging in, the user's IP address. Figure 7 illustrates a user's Talk page.

Figure 7

The screenshot shows a Wikimedia Meta-Wiki user talk page. At the top, there's a navigation bar with 'User page' and 'Discussion' tabs. Below that is a search bar and a language selector set to 'English'. The main heading is 'User talk:Jalexander-WMF'. A table of contents lists 13 items, starting with 'Welcome to Meta!'. The main content area contains a 'Welcome to Meta!' section followed by a discussion thread. The thread title is 'About the "Notification of Wikimedia Foundation actions regarding local CheckUser" in Chinese Wikipedia'. The first post is by user fr33kman, dated 05:21, 17 August 2010 (UTC). It says 'Congrats on the job'. A reply from Jalexander-WMF, dated 01:18, 19 August 2010 (UTC), says 'I am sorry to interrupt you with poor English to talk with you here. I hope WMF and you know that that is Chinese Wikipedia, that we may not have good English to request for User Check. I just want to ask if we may use Chinese to request for it? Actually I knew English, but not much.' A reply from Sanmosa, dated 07:37, 30 March 2018 (UTC), says '@Sanmosa: While I imagine that it will go a bit faster if requested in English (or another language the Stewards already understand) I am sure they will be ok with you asking in any language and I will certainly try to help them however possible with translation etc. Jalexander-WMF 05:24, 5 April 2018 (UTC)'. Another reply from Jalexander-WMF, dated 06:02, 5 April 2018 (UTC), says 'Thank you for your reply. Anyway, I have a great news to you, that the Chinese Wikipedia Community had set up a translation page of CheckUser request from Chinese to English, to reduce misunderstanding, so that Stewards may read the requests directly. However, actually, I would use English to request for CU directly, with my hopefully average English.' A final reply from Sanmosa, dated 05:30, 5 April 2018 (UTC), says '@Sanmosa: That's great to hear and a good idea! I do think you should be ok asking directly in English if desired, many (most) of the Stewards are also native to different languages and are used to working with those who speak less English. Jalexander-WMF 06:02, 5 April 2018 (UTC)'. A final reply from Jalexander-WMF, dated 06:08, 5 April 2018 (UTC), says 'OK, I just try to use more English, to broaden my horizons.'

24. Users may also interact with one another and with Wikimedia using private or semi-private communication features on Project websites. The primary examples of these features are discussed below.

- **OTRS.** The Open-source Ticket Request System is a software package employed by Wikimedia since 2004 to address queries, complaints, and comments about the Wikimedia Projects from users. Responses are handled primarily by a group of trusted Wikimedia users who volunteer to answer these queries, complaints, and comments, and certain responses are handled by Wikimedia staff. In other words, OTRS messages are private communications between Wikimedia users or between users and Wikimedia staff.
- **Email User Feature.** This feature allows registered users to exchange emails with one another via Wikimedia, provided that both users have enabled email communications on their Wikimedia accounts. For example, a Wikimedia user can send an email to another Wikimedia user's Gmail account by visiting the intended recipient's User page and selecting an email option. The content of the communication is entered into an HTML form and transferred via HTTPS to Wikimedia servers, and from there to the recipient's email address using the standard SMTP email protocol.

- *Private Mailing Lists.* Users may create or be invited to use “private mailing lists” through Wikimedia. Users must seek permission from the mailing list administrator to receive access to, send, or receive emails through the list, and to read historical archives of mail sent to the list. List messages are sent and received using the SMTP protocol, while archives are accessed via HTTP/S (both via Wikimedia’s servers at lists.wikimedia.org). For example, Wikimedia Israel Galleries, Libraries, Archives and Museums coordination is one such private mailing list: <https://lists.wikimedia.org/mailman/listinfo/wikimedia-il-glam>.
- *Private Wikis.* Some groups of users, for example OTRS volunteers, also employ “private wikis” to communicate about specific topics. These wikis are only accessible by certain users. Wikimedia servers host these communications, which are transmitted via HTTP/S.

### III. STATISTICS ON WIKIMEDIA’S INTERNATIONAL INTERNET COMMUNICATIONS

25. As the operator of one of the most-visited websites in the world, Wikimedia engages in an extraordinarily high volume of Internet communications with individual users located around the globe, some of whom are Americans located abroad. *See, e.g.*, Exhibit 10 (WIKI0009221, -9222) (American citizen and Wikimedia user located abroad requesting ability to use web application designed to anonymize online activity). Indeed, its servers handle more than one trillion international communications each year, with individuals who are located in every country on earth.

26. Wikimedia’s international communications can be divided into the following three categories, each of which contain sensitive and private information related to Wikimedia and its users:

- **Category 1 – *Wikimedia communications with its community members.*** Examples of these communications include, but are not limited to, page views to Wikimedia websites, edits and contributions to Wikimedia websites, emails between registered Wikimedia users and emails on Wikimedia’s mailing lists. The vast majority of these communications are generated by HTTP and HTTPS requests from users who read and contribute to Wikimedia’s Projects, and who use the Projects and webpages to interact with each other.
- **Category 2 – *Wikimedia’s internal log communications.*** Wikimedia maintains proprietary “logs” that catalogue a variety of information regarding each HTTP or HTTPS request to a Wikimedia resource. These logs help Wikimedia monitor, study, and improve the Projects. Every time Wikimedia receives an HTTP or HTTPS request from a person

accessing a Project webpage, it creates a corresponding log entry. Depending on the location of the user and the routing of that user's request, the log may be generated by Wikimedia's servers abroad, which in turn send the log entry to Wikimedia in the United States. Wikimedia also uses information about the log communications to inform its site operations, improve user interfaces, and to guide support work for its volunteer community.

- **Category 3 – Wikimedia's staff communications.** Wikimedia's office network router located in the United States also handles a variety of border-crossing communications. Examples of these communications include, but are not limited to, Gmail, Google chat, Internet Relay Chat, and Slack. Additionally, Wikimedia staff members use a variety of third-party tools to conduct their work, including, but not limited to, Google Apps/G Suite, Trello, Sugar, Qualtrics, UserTesting and Salesforce. Wikimedia's staff communications reveal a variety of highly sensitive information about Wikimedia's operations, and include communications with international staff who are involved in political or social activism and live or work in geopolitical areas that are a special focus of the U.S. government's counterterrorism or diplomatic efforts.

27. Wikimedia is able to use its server log data to quantify each of these categories of international communications by Internet protocol and country of origin. I led the process of calculating these communications volume figures in response to Defendants' written discovery requests. The results of these analyses are contained in Wikimedia's Exhibit 14 and described below:

**Category 1 Volume – Wikimedia communications with its community members**

- Total HTTP and HTTPS requests from foreign users to Wikimedia's U.S. servers: Between August 1, 2017 and January 31, 2018, Wikimedia's U.S. servers received approximately **381 billion HTTP/S requests** from users outside of the United States. Wikimedia calculated this figure using MaxMind<sup>2</sup> geolocation data to determine the country associated with the client IP of each HTTP/S request transmitted to Wikimedia's servers in the United States.
- Total HTTP and HTTPS requests from U.S. users to Wikimedia's foreign servers: Between August 1, 2017 and January 31, 2018, Wikimedia's foreign-based servers received approximately **2.8 billion HTTP/S requests** from users in the United States. Wikimedia calculated this figure using MaxMind geolocation data to determine the country associated with the client IP of each HTTP/S request transmitted to Wikimedia's servers located outside the United States.

---

<sup>2</sup> MaxMind is a U.S.-based company that provides widely used databases for Internet geolocation.

**Category 2 Volume – *Wikimedia’s internal log communications***

- Total log communications transmitted from Wikimedia’s foreign servers to Wikimedia’s U.S. servers: Between August 1, 2017 and January 31, 2018, Wikimedia’s foreign-based servers transmitted approximately **736 billion log communications** to Wikimedia servers in the United States, using the Apache Kafka protocol. These communications were encrypted using IPsec, a secure network protocol suite that authenticates and encrypts the packets of data sent over a network.

**Category 3 Volume – *Communications by Wikimedia staff***

- Total logged international TCP connections using Wikimedia’s office network or VPN: Between March 1, 2017 and February 28, 2018, Wikimedia’s office network router located in the United States logged open Internet connections using the Transmission Control Protocol (TCP) with non-U.S. countries, territories and regions at least approximately **4,948,011** times.
- Total logged international UDP connections using Wikimedia’s office network or VPN: Between March 1, 2017 and February 28, 2018, Wikimedia’s office network router located in the United States logged open Internet connections using the User Datagram Protocol (UDP) with non-U.S. countries, territories and regions at least approximately **2,207,771** times.
- Total logged international ICMP connections using Wikimedia’s office network or VPN: Between March 1, 2017 and February 28, 2018, Wikimedia’s office network router located in the United States logged open Internet connections using the Internet Control Message Protocol (ICMP) with non-U.S. countries, territories and regions at least approximately **51,301** times.

28. Wikimedia calculated the Category 3 communications using MaxMind geolocation data to determine the country associated with the source and destination IPs of each of these Internet connections. The router generated a log of all open Internet connections, including their protocol and source and destination IPs, every five minutes, using a program called “contrack.” Connections that remained open for more than five minutes may have been logged more than one time. Connections that remained open for less than five minutes may not be logged at all. These figures include, but are not limited to, connections sent through Wikimedia’s virtual private network (VPN), which passed through the same router and can be distinguished in the log by the IP address used to connect.

29. In response to discovery requests in this litigation, I also quantified and verified other statistics relevant to the volume of Wikimedia’s communications, a subset of which are contained below in Figure 8.

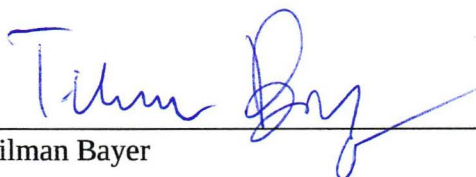
**Figure 8**

*Wikimedia Statistics and Volume Metrics*

	<b>Number</b>	<b>Date Range</b>
<i>Wikipedia articles</i>	47,433,067	As of February 9, 2018
<i>Wikipedia languages</i>	288	As of February 9, 2018
<i>Edits over the lifespan of the Wikimedia Projects</i>	approx. 3.4 billion	As of February 28, 2018
<i>Range of monthly unique devices visiting Wikipedia pages</i>	1,521,315,531 (high) / 1,361,541,637 (low)	High/low range for April 2017 to December 2017
<i>Range of monthly U.S. unique devices visiting Wikipedia pages</i>	284,472,345 (high) / 261,065,481 (low)	High/low range for April 2017 to December 2017
<i>Page views</i>	237,586,909,120	January 1, 2017 to December 31, 2017 (one year)
<i>U.S. page views</i>	74,182,143,412	January 1, 2017 to December 31, 2017 (one year)

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Executed on December 19, 2018 in Sindelfingen, Germany.

  
 \_\_\_\_\_  
 Tilman Bayer